RNR-88-014

NASA-TM-112078

IN-61-TM

013406

- **Means to test technology**

A consortium member can evaluate supercomputer performance before opting for single ownership.

- **Local economic benefits.**

A community's ability to offer local researchers access to a consortium-owned supercomputer can help retain high-level expertise in the area.

Consortium: The Cons

- **Security risk**

When multiple parties share the same supercomputer, steps must be taken to protect the transmission and storage of sensitive data.

- **Conflict over resources**

Because processing loads can rarely be projected accurately, consortium members can end up contending for system resources.

- **Partnership considerations**

As with any partnership, in a consortium conflicts can also arise over which resources are needed most and how many dollars should be spent on those resources.

Time sharing: The Pros

- **No capital investment**

Time-share users avoid the costs of purchasing, installing, maintaining, and running supercomputer hardware and software. In addition, today's time-share services include a range of flexible pricing options aimed at keeping users' computing budgets intact.

- **Application libraries**

As part of a time-share service, users can access the service provider's application library and avoid purchasing costly software licenses for applications that will be used only occasionally.

- **Responsive support**

Time-share providers compete largely on their ability to respond to users' price/performance/job-turnaround requirements.

- **Low cost during low activity**

Because time-share users incur none of the ongoing expenses associated with operating and maintaining a supercomputer inhouse, supercomputing costs plummet when processing requirements are at a minimum.

- **Growth path to own**

Time-share users can evaluate supercomputer performance and obtain the needed processing cycles until utilization increases to the point that it becomes more cost-effective to bring a supercomputer inhouse.

Time sharing: The Cons

- **External resource**

Time-share users have less control over their computing resources than do companies that bring a supercomputer inhouse.

- **Greater accountability**

The use of a remote supercomputing service incurs greater budget accountability than does use of an internal resource.

- **High cost during high activity**

During periods of heavy usage, it can be more expensive to access a supercomputer remotely than to use an installed system. ■

REPORT:

SPARSE MATRIX ALGORITHMS WORKSHOP

HORST D. SIMON

It has long been recognized that the efficient solution of sparse linear equations is a key element for solving large scientific applications problems. A recent study lists several major scientific application areas which require at some point the numerical solution of linear systems: lattice-gauge theory, structural analysis, circuit simulation, device simulation, quantum chemistry, computational fluid dynamics, and geodetic networks. The availability of supercomputers and breakthroughs in algorithm research have led to the feasibility of larger and more refined models in these scientific applications areas, yet new and more demanding models remain to be solved. It is clear to most researchers involved in these fields that only the combined efforts of advanced algorithm technology and new supercomputing hardware will enable us to solve the challenging problems ahead. However, quite often scientists involved with supercomputing applications are unaware of the latest algorithmic advances, and algorithm researchers frequently perform their work unaware of pressing applications needs. This situation has been recognized, and the SIAM-sponsored "National Computing Initiative" calls for support of computational science and engineering research in an interdisciplinary mode.

In order to address some of these challenging problems, and in order to make a step towards interdisciplinary cooperation, NASA Ames Research Center, Boeing Computer Services, and Cray Research, Inc., sponsored a three-day workshop in March. The workshop was organized by Phuong Vu of Cray Research and the author. The goal of the workshop was to bring together all three groups involved in the progress of supercomputing: hardware designers, algorithm developers, and researchers in scientific applications. The goal of the workshop was to provide a better interaction between the various groups involved in supercomputing. The numerical solution of sparse linear systems was used as the unifying theme of the workshop.

In invited presentations, algorithm researchers provided a survey of the current state of sparse matrix research on supercomputers. Direct sparse matrix methods were discussed by Iain Duff from AERE, Harwell, England, and Barry Peyton from Boeing Computer Services in Seattle. Peyton dispelled the myth that general sparse methods would not vectorize by showing results of 140 Mflops and higher on a Cray X-MP. Youcef Saad (University of Illinois and RIACS) presented a survey of iterative methods, and Henk van der Vorst (Delft University of Technology, The Netherlands) showed very impressive results on how to speed up incomplete factorization techniques. All algorithm talks demonstrated that significant progress had been made in developing powerful tools for solving linear systems.

The applications speakers came from very diverse areas. Gary Stanley and Bahram Nour-Omid from Lockheed Palo Alto Research reported on computational problems in structural mechanics. In a research project supported by NASA Langley, they are looking towards routinely solving systems of the order of 100,000 to 500,000 degrees of freedom. Kurt Marfurt from AMOCO Oil in Tulsa is currently solving problems with a million degrees of freedom on a Cray-2 in production runs in seismic applications. He does not yet know how to do three-dimensional problems. Richard Smith from Exxon gave an overview on the current requirements in reservoir modeling, and a group from Boeing (Larry Wigton and John Bussoletti) presented an efficient combination of direct and iterative methods, which allows them to solve very large three-dimensional problems. Other applications talks covered such areas as semiconductor device simulation (Tony Shober, AT&T Bell Laboratories), linear programming (Kathryn Turner, Utah State), and chemical engineering (Mark Stadtherr, Univ. of Illinois). Al Erisman from Boeing gave some historical highlights of sparse matrix methods, and Steve Nelson from Cray Research summarized technological trends and their impact on the Cray Y-MP line of supercomputers.

The discussion sessions distinguished the workshop markedly from the usual conference. Algorithm researchers and applications researchers met to address open research questions, define hardware requirements, and to exchange ideas for the solution of some of these difficult problems. The software discussion group was involved with issues such as the migration to a parallel environment, the use of general-purpose packages in a parallel environment, and software tools for scientific computing. Other groups discussed unresolved algorithmic issues, the use of advanced parallel architectures, and questions of standardization.

Participants in the workshop stated that they derived many benefits from attending the workshop. Tony Shober from AT&T Bell Laboratories, whose group is developing a sparse matrix package from device simulation, expressed his experience with the workshop as follows: "We had a great opportunity to talk with all the experts and found that we do to 70 percent exactly the right thing. To have this affirmation is very good. For the remaining 30 percent we are bringing back several new ideas, and this is even better."

Scientific applications users of supercomputers now have a better overview of recent progress in algorithm research for sparse linear systems. They learned about new techniques developed for different disciplines and are able to leverage new methods directly into their application codes. Participation in the workshop will hopefully shortcut the often lengthy process of disseminating new algorithms from the applied mathematics and computer science community into the application areas (and vice versa). Some algorithm researchers have developed a better understanding of the computational requirements of current and future applications. They learned about challenging new problems pursued in the applications, which wait for a new algorithmic approach.

Participation in the workshop will open up for them new areas of research with immediate practical applications. Hardware designers developed a better understanding of the computational requirements of current and future applications. It is hoped that supercomputer manufacturers will integrate many of the suggestions of the workshop participants into future hardware and software products. ■

Dr. Horst D. Simon is a senior researcher with Boeing Computer Services. He received his undergraduate degree in mathematics from the Technical University Berlin, Germany, in 1978, and his Ph.D. in Mathematics from the University of California, Berkeley, in 1982.

He joined Boeing Computer Services in 1983 as a research analyst. His assignments included the development of mathematical software for the CRAY 1-S and CRAY X-MP computers.